

04/25/00
jc811 U.S. PTO

4-26-00

A

IBM Docket No. RSW9-2000-0039-US1

**In the United States Patent and Trademark Office
Patent Application Transmittal**

Transmitted herewith for filing is the Patent Application of:

Inventors(s): Christopher A.S. Gage, John R. Hind, and Marcia L. Peters

For: URL-Based Sticky Routing Tokens Using a Server-Side Cookie Jar

jc530 U.S. PTO
09/557708
04/25/00

Enclosed are

22 pages of specification, including 27 claims, plus 7 sheets of *formal* drawings.

An assignment of the invention to International Business Machines Corporation, Armonk, New York 10504.

A certified copy of a/an application.

X Declaration and Power of Attorney. Not executed
PTO-1449 & references

X A return post card

X Other: Memo re Missing Parts

Filing Fee Calculation (For Other Than Small Entity)

Basic Fee:							\$690.00
Claims Fees:		Filed	Limit	Extra		Rate per Extra	
Total claims:		27	20	7		\$18.00	\$126.00
Independent claims:		7	3	4		\$78.00	\$312.00
1	Multiple Dependent Claim Presented					\$260.00	\$260.00
Total							\$1,388.00

Please charge Deposit Account **09-0461** for the **Total** set forth above. The Commissioner is authorized to charge payment of any additional filing fees required under 37 CFR §1.16 and any patent application processing fees under 37 CFR §1.17 or to credit any overpayment to the identified account. A duplicate copy of this sheet is enclosed.

Express Mail Certificate

Express Mail Label No: EJ922406387US

Date: April 25, 2000

I hereby certify that I am depositing the papers identified above with the U.S. Postal Service "Express Mail Post Office to Address" service on the above date, addressed to the Commissioner of Patents and Trademarks, Washington, DC 20231

Dianne Lane

Dianne Lane

BY: *Jeanine S. Ray-Yarletts*
Jeanine S. Ray-Yarletts
Attorney of Record Reg. No. 39,808

Date: April 25, 2000

IBM Corporation T81/062
Intellectual Property Law
PO Box 12195
Res. Tri. Park, NC 27709

Telephone: 919-543-2541 FAX 919-254-4330

EXPRESS MAIL LABEL NO. EJ922406387US

DATE OF DEPOSIT: April 25, 2000

I hereby certify that this paper and fee are being deposited with the United States Postal Service Express Mail Post Office to Addressee service under 37 CFR §1.10 on the date indicated above and is addressed to the Assistant Commissioner of Patents, Washington, D.C. 20231.

Dianne Lane

Dianne Lane

NAME OF PERSON MAILING PAPER AND FEE

SIGNATURE OF PERSON MAILING PAPER AND FEE

URL-Based Sticky Routing Tokens Using a Server- Side Cookie Jar

Background of the Invention

Through the 1990s, computer networks have grown exponentially. The Internet and the worldwide web have allowed everyone access to extraordinary amounts of information. They have also allowed people to conduct much of their day-to-day business, such as shopping or research using the Internet. During the 1999 shopping season, several large department and toy stores did more business while their stores were actually closed than when they were opened; this was due to their presence on the Internet. Their Internet web sites allowed their customers to look at the products they carry, search for particular items, read about the items, place orders and view delivery information all after the children were tucked quietly in bed. This came as quite a surprise to a large portion of the industry. Merchants had to increase the size and number of their servers to support all of the traffic from the home users. Several had to install significant amounts of new hardware such as network dispatchers to route

requests to their servers to redundant or back-up servers so that the response time to the users did not become unbearable.

This is just one of the many examples of how the Internet is growing and how the access to information over the Internet is impacting society today. Users are becoming accustomed to receiving almost instantaneous response from computer networks and will not tolerate delays. As the number of users of the Internet grows, the number of redundant servers and network dispatchers managing these servers also must increase to maintain the quality of service that the customers demand. Network dispatchers are used to manage requests to a server, and distribute the work among several redundant servers, using a predetermined load-balancing methodology.

A typical network dispatcher system is shown in Figure 6. When the end user device **601** requests information from a server that has a front-end network dispatcher **603**, the network dispatcher receives the request and routes it to one of the redundant back-end servers **605**, **607**, **609**. If the network dispatcher **603** subsequently receives another request from the end user device **601**, it will go through its selection process again to maintain the continuous load balancing. This is the function of a network dispatcher as it was designed.

The use of the traditional network dispatcher can become a problem when a user, once routed to a server for a first request, must consistently be routed to that same server for repeat requests due to the collection of information. This is so, in the merchant case, where a user comes into the merchant's web site and begins to place an order. Once the first item is placed into a shopping basket, the items must be remembered so that the user can continue to shop and place a successful order.

More specifically, suppose that a cluster of web servers provide equivalent services, front-ended by a network load balancer. The load balancer's job is to route inbound packets to the least-busy server, based on decision-making mechanisms that are

beyond the scope of this discussion. A simplified version of this is shown in Figure 1 where a cellular phone **101** is accessing information on redundant servers **109** using a wireless link **115** to a cellular tower **105**, then a land link **117** to a dispatcher **107**, then to the information servers **109**, while a notebook computer **103** is also accessing information on the same servers **109**. The problem becomes how to load balance given that a client, such as the cellular phone **101**, must return to a particular server in the cluster **109** for all the flows comprising a session or unit of work. This will be referred to as the "sticky routing" problem.

In the present document, a "session" is defined as a series of related transactions to perform a unit of work. A session generally utilizes HTTP (hyper-text transport protocol) or HTTPS (secure hyper-text transport protocol) flows, consisting of one or more TCP/IP (transmission control protocol/Internet protocol) connections. A simple electronic commerce transaction typically consists of a sequence of related actions such as browsing an online catalog, selecting one or more items of merchandise, placing the order, providing payment and shipping information, and finally confirming or canceling the entire transaction. Information about the state of the session may span multiple TCP/IP connections, since information such as the client's identity, the item desired, the agreed-upon price, payment information, etc. must persist until the entire transaction is complete.

When a given client has a session with a particular server, state information about that session exists only at that particular server. In this case, a load balancer needs to apply extra intelligence to route the packets correctly. In particular, it needs to choose the same server repeatedly as the destination for all inbound packets from a given client for a given session or transaction. This client-server relationship will be referred to as "binding". To load-balance effectively over time, the system must also release the client's affinity for a particular server between sessions or transactions.

mobile clients to the Internet, such as a Wireless Access Protocol (WAP) gateway, and also appears in home networking devices such as a LAN router or smart hub or in modems for the home (3Comm's ISDN LAN modem is an example of a small router for the home incorporating NAT function). NAT devices and transparent proxies are also
5 deployed by ISPs offering "always on" types of services such as those based on cable modem or Asymmetric Digital Subscriber Loop (ADSL) technology as well as in traditional dial-up "Point of Presence" (POPs).

The SSL-ID (Secure Sockets Layer Identifier) has also been tried as a solution to the sticky routing problem and failed. Connections using the SSL or TLS are encrypted.
10 Once an SSL connection is established between a given client and a particular server, the SSL ID (a quasi-unique number) could be examined by the load-balancer and used for sticky routing purposes. Although the SSL standard always permitted either endpoint of the connection to repudiate the key agreement and force a renegotiation of SSL parameters and consequently the assignment of a new SSL-ID, in practice only
15 servers did this, making the approach viable for a while. However, with the recent release of Microsoft Internet Explorer 5.0, this technique is no longer viable. Internet Explorer 5.0 is coded such that either the server or the client may repudiate the key agreement, making it impossible for a load-balancer to correlate the former SSL connection with the current one.

The next solution attempted for both of these problems employed "cookies". A cookie
20 is a data object transported in variable-length fields within the HTTP header that is normally stored on the client, either for the duration of the session or permanently. A cookie stores certain data that the server application wants to remember about a particular client. This could include client identification, session parameters, user
25 preferences, session state, or almost anything else an application writer can think of. Although a load-balancer with content-based routing could look into the HTTP header and route based on data contained in cookies, this initially promising solution also turned out to have a disastrous flaw. Certain clients are incapable of storing cookies.

These certain clients include webphone clients that access the Internet through a WAP gateway using the Wireless Session Protocol (WSP). WSP does not include cookies. Even if WSP supported cookies, the webphone clients are not capable of storing cookies due to their extremely limited memory. While it is possible for a wireless gateway product to store cookies on behalf of the wireless client (the IBM eNetwork Wireless Gateway does this; the Nokia WAP gateway does not), such functions in the gateway cannot be assumed, as is demonstrated above. In addition, with increasing privacy concerns about the use of cookies by unscrupulous advertisers to track an Internet user's surfing habits, many users are choosing to disable cookies altogether, or turn on cookie prompting, accepting cookies selectively, if at all. So the capability of storing persistent session information in cookies cannot be presumed.

Summary of the Invention

The present invention enables sticky session persistence to a given server for individual clients, even if they have identical IP addresses, such as WAP phones accessing the network through a NAT gateway. It also restores web applications' ability to rely on the presence of cookies, which they had lost with the advent of "cookie-free" WAP phones. The present invention provides better support for robust, large-scale, reliable, highly-available electronic commerce installations serving all types of web clients, including Microsoft Internet Explorer 5.0 and the newer webphone and mobile devices that lack cookie support. The present solution can be delivered as a server platform-based service, without requiring any changes in existing web applications.

The present invention utilizes a modification to the Uniform Resource Locators (URLs) in a Hypertext Transport Protocol (HTTP) or Secure HTTP (HTTPS) document such that the URL uniquely identifies a given client and binds it to a particular server for the duration of a session. In addition to the modification to the URL, a "server-side cookie jar" is used which provides a data object accessible to a web application server. The

cookie jar is used to store cookies on behalf of the particular client so that the clients or their client-side proxies will not be required to store them. Using the modified client-unique URL to identify a specific server, client, session and a cookie jar in an inbound web request, the request is routed to the appropriate server. The present invention will be described in greater detail with respect to a preferred embodiment below.

Objects of the Invention

It is an object of the present invention to enable users of the Internet, unable or unwilling to store local cookies, or located behind a NAT or transparent proxy, to establish a session with a particular server in a cluster of servers front-ended by a dispatcher.

It is another object of the present invention to enable the server, in a set of dispatcher managed servers, to store information regarding the client.

It is a further object of the present invention to enable this session to occur by modifying the URL and requiring no unique modifications of the clients.

It is yet another object of the present invention to enable an entire transaction to occur between a client and one of a plurality of servers managed by a network dispatcher.

It is still a further object of the present invention to enable differentiated quality of service using the described methodology.

Brief Description of the Drawings

Figure 1 is a pictorial representation of a minimal network in which the present invention can operate.

Figure 2 is a diagram of the changes to the URL for the present invention.

Figure 3 is a flow diagram for the changes to the URL in the present invention.

Figure 4A depicts the logic flow for the receipt of information to the network dispatcher of the present invention.

5 Figure 4B depicts the logic flow for the receipt of information from the application.

Figure 5 demonstrates the use of Network Address Translators (NATs).

Figure 6 shows how depicts how a typical network dispatcher would route information.

Detailed Description of the Preferred Embodiment

10 The above mentioned objectives, as well as others, will be described in detail with respect to a preferred embodiment of the present invention as well as to the Figures presented herein. Like numbers in the figures represent the same elements. The preferred embodiment is presented as an example only and is not meant to limit the stated invention or claims in any manner.

15 The preferred embodiment of the present invention utilizes a modification to the URL in an HTTP or HTTPS document such that the document uniquely identifies a given client and binds the client to a particular server for the duration of a session. In conjunction with the URL modification, a server-side cookie jar is implemented that enables the server to store cookies on behalf of a particular client. This releases the client or client-side proxies from the responsibility of storing the cookies. The cookie-jar approach
20 renders the issue of cookie-storage on the client moot. Using the modified client-unique URL to identify a specific server, client, session and cookie jar in an inbound

web request, the request is routed to the appropriate server where an inbound data stream filter restores the URL to its unmodified state, retrieves the appropriate cookie or cookies from the designated cookie-jar, and inserts them into the HTTP stream before passing the inbound request to the application (or to the next inbound data stream filter in the chain, if filter chaining is used). The inbound data stream filter must be the first or only filter to operate on the inbound data stream for the present invention to successfully accomplish its objective. A paired outbound filter, also resident on the web application server, receives the application-generated outbound data and headers, moves any cookies found in the headers into a specific cookie jar associated with the client's session, and converts certain URLs in the outbound web page to the modified form that will be described in detail below. The outbound data stream filter must be the last or only filter operating upon the outbound data stream.

The outbound data stream filter of the preferred embodiment inserts a string delineated with a forward slash character, the string preferably being formatted using modified Base64 encoding (the special characters "+" and "/" in the standard encoding would be replaced by "-" and "_" respectively), into the URL between the server portion and the path portion. Modified Base64 encoding is preferred so that arbitrary binary data in the inserted portion is represented as a series of legal URL path printable characters, although any means of accomplishing this is acceptable. An example of this is shown in Figure 2. Figure 2 shows the unmodified URL **201** as well as the modified URL **203**. The string, hereafter called a "sticky routing token" **235**, contains four fields:

1. A routing field identifying the specific server **207** to which the client session is bound. Preferably this routing field is the server's IP subnet address on the same IP subnet as the load-balancer, although other means of server identification are acceptable.
2. A date and time stamp **209**. A fresh date-time stamp is inserted by the outbound filter into the outbound data stream. It is also inspected by the load balancer and the

inbound data stream filter to determine if the binding relationship between the client and the server is stale.

3. A "key" **211** or index that can be used to select the proper cookie jar where cookies pertaining to the specific client-server binding are kept.

5 4. A checksum or hash **213** (such as a SHA-1 Secure Hash Algorithm) over fields 1-3 that can be checked to distinguish a valid sticky-routing URL from an unmodified URL.

The outbound filter registers to handle only certain MIME (Multi-Purpose Internet Mail) types (documents encoded using a structured markup language including but not limited to HTML, WML and XML). It does not register to handle "streaming media", image, or downloaded code MIME types.

Referring now to Figure 4A, before a client-server binding is created, the load-balancer will see only a standard unmodified URL in the inbound data stream. The load balancer uses standard load- balancing technology, outside the scope of this discussion, to route this initial inbound request to the most appropriate server **401**.

When this request arrives at the application server and no binding exists, the inbound data stream filter of the present invention creates a sticky routing token **403** for later insertion into the outbound data stream, as described above. The third field of the string is a new key that accesses a storage object in which cookies for the session can be stored **405**. This key also serves to correlate the inbound and outbound streams associated with a given binding. In one preferred embodiment, the inbound filter creates a special cookie into which it inserts the just-created key and places this cookie into the header. This "key cookie" is used to store state information, namely the key itself, during the data passage from the inbound filter, to the application, and back to the outbound filter. Alternatively, the inbound filter can store the key in a connection control

object associated with the TCP connection. Although exactly where the key is stored is not of particular importance to implementing the present invention, it is important to store the key in a place where it can later be retrieved by the outbound data stream filter in a manner that enables session correlation.

- 5 The inbound filter then forwards the data to the next filter in the chain **409** if there are multiple inbound filters **407**, or directly the application **411** if there is only one inbound filter.

10 Referring now to Figure 4B, after the application has processed the inbound data and created an outbound structured markup language stream, the outbound data stream filter of the present invention gains control **421**. At this point the header may contain cookies in which the application has stored state information pertaining to the specific client. If the inbound filter had used the "key cookie" technique **423** described in the present invention to pass the key to the outbound filter, this key cookie will also be present in the header. The outbound filter retrieves the key for the session **425**,
15 removes all cookies **427** from the header, and stores them **429** in the cookie jar indexed by the key. It is important to note that outbound filter of the present invention could be implemented as part of a transcoding process, and should not significantly increase outbound data stream filtering path length if the data is already being parsed, formatted, and copied on the outbound data path during a transcoding operation.

- 20 Next, the outbound filter modifies zero or more of the URLs within the structured markup language document **431**. All URLs referring to the same transaction need to be modified. In practice, all URLs referring to the particular server, or relative to that server, are modified. (Alternatively, a software-development tool could be provided enabling the application programmer to mark certain URLs as related to the transaction,
25 in which case only those would be modified by the outbound filter; however this implementation is less preferred, due to its requirement that application programs be modified.) The outbound filter of the present invention creates the sticky routing string

433 mentioned above, and freshens the date-time stamp 435. The modification consists of inserting this sticky routing string into each of the selected URLs 437 (this process is sometimes called "URL-rewriting"). In the preferred embodiment, this might easily be achieved by invoking an existing URL-rewriting function with new parameters. (For additional information on the state of the art for URL-rewriting, see IBM WebSphere Application Server documentation at [http:// web/doc/whatis/icesessta.html](http://web/doc/whatis/icesessta.html) . This document discusses the current state of the art in session-state correlation, cookie management, and URL-rewriting.) Finally the output filter of the present invention passes the resulting stream to the network layer for transmission to the client 439.

When the client receives the modified HTTP stream, the user may select any one of the modified URLs, resulting in an inbound request containing a sticky routing token. The present invention is advantageous over prior art techniques that exploited optional or variable features of the protocol such as the source IP address, SSL-ID, or cookies, since the URL is mandatory HTTP content that cannot be removed, disguised, or made optional.

Referring once again to Figure 3, the inbound request will arrive at a load-balancer 301. The load-balancer inspects the data to determine if a valid sticky routing token token is present 303. In the preferred embodiment, validation consists of comparing the embedded checksum or digital signature (hash, such as a SHA-1 hash) with a computed checksum or digital signature. If a valid sticky routing token is not present, it handles the packet as usual 305. The load balancer then tests the date-time stamp to determine if the session binding is stale. If the difference between date-time stamp and the present date-time exceeds some constant, the binding is considered stale and the data is processed as if no sticky routing field were present, i.e. preferably by applying existing load-balancing techniques to select a server to handle the request then creating a new session binding and a corresponding sticky routing field. In either event, the routing field in the token is used to route the packet to the identified server 313.

Upon the packet's arrival at the application server, the inbound data stream filter gains control.

The inbound data stream filter validates the sticky routing token and date-time stamp in the same manner explained above. If the sticky routing token is valid and not stale, the inbound filter uses the key to access the cookie-jar where cookies for the particular client-server binding are stored. It removes the routing token and saves the key either in a key cookie or in the TCP connection control block, as described above. It inserts all the selected cookies into the header and forwards the data to either the next inbound filter, if any, or directly to the application.

The sticky routing token can also serve as the basis for a load-balancer to provide differentiated quality-of-service. By being able to recognize one particular client session or distinguish a transaction-in-process from one that has not yet started, appropriate decisions can be made to prioritize certain packets over others.

The present idea can also be extended so as to provide a high-availability implementation without substantially changing the underlying theory. To extent this, instead of a cookie-jar or object store per application server, instead there is a shared low-overhead Object Store accessible to all servers in the cluster (and, optionally, to the load-balancer).

Instead of storing the cookies in a cookie jar, cookies are stored in the shared object store. In addition to cookies, the application would be rewritten to store all of its vital state information concerning an ongoing session into the object store. The application would need to extract all of its state information from the object store upon receipt of any inbound data.

What is claimed is:

1 /- A method of establishing a persistent relationship between an end user device
2 and a server where the server is one of a plurality of servers managed by a dispatcher
3 and the end user device accesses the server using a universal resource locator (URL),
4 the method comprising the steps of:
5 receiving, at the dispatcher, a request for information from the end user device;
6 determining, by the dispatcher, which of the plurality of servers to select for
7 satisfying the request;
8 creating, at the selected server, a token comprising at least an identifier for the
9 selected server, a date/time stamp, and a key, said key for accessing a storage area for
10 information regarding the persistent relationship and the end user device;
11 inserting the token into the URL; and,
12 sending, by the selected server to the client device, a response with the token
13 inserted into the URL.

1 2. A method as claimed in claim 1 wherein said token is encoded using a modified
2 Base64 encoding.

1 3 - A method as claimed in claim 1 wherein said token has a checksum or hash
2 verification field.

1 4. A method as claimed in claim 3 wherein said hash is a SHA-1 hash computed
2 over said identifier for said selected server, said date/time stamp, and said key.

1 5. A method as claimed in claim 3 wherein said checksum or hash is encoded
2 using a modified Base64 encoding.

1 6 - A method as claimed in claim 1 wherein said information regarding the persistent
2 relationship is stored as a cookie on said server.

1 7 - A method of routing a request by an end user device to a particular one of a
2 plurality of redundant servers residing behind a network dispatching mechanism, said
3 methods comprising the steps of:
4 receiving, at the network dispatching mechanism, a request for information
5 indicated by a uniform resource locator (URL);
6 determining, at the network dispatching mechanism, if said URL contains a valid
7 routing token;
8 if said URL contains a valid routing token, further determining, at the network
9 dispatching mechanism, if a session binding indicated by said routing token is old;
10 if said URL contains a valid routing token and said routing token is not old,
11 forwarding, by said network dispatching mechanism, the request, including the URL, to
12 the particular server indicated by said valid routing token;
13 removing, by said particular server, said valid routing information from the URL;
14 storing, by said particular server, routing information removed from said valid
15 routing token, where said valid routing information can be accessed subsequently by an
16 outbound data stream filter during the processing of an outbound reply related to said
17 request;
18 accessing, by said particular server, a storage location where information
19 regarding a session between the particular server and the end user device is stored;
20 and
21 inserting, by said particular server, said session information into said request.

1 8 - The method as claimed in claim 7 wherein additional filtering of the URL is done
2 prior to the forwarding step.

1 9- The methods claimed in claim 1 thru 8 wherein all filtering is performed within the
2 dispatcher.

1 10- A method of sending information to a requesting end user from an application
2 over a session wherein said application resides at one of a plurality of redundant
3 servers residing behind a network dispatcher, said method comprising the steps of:
4 receiving response information from said application, said response information
5 including a URL (uniform resource locator);
6 determining if a key cookie has been used for storing session information
7 between said end user and said application;
8 if a key cookie has been used for storing session information, retrieving a
9 session key from said key cookie;
10 if a key cookie was not used for storing session information, retrieving said
11 session key from a control block;
12 removing all cookies from said response information;
13 storing said removed cookies in a predetermined storage area;
14 updating said URL to indicate the removal of said cookies;
15 creating a sticky routing string;
16 updating a date/time stamp in said sticky routing string;
17 inserting said sticky routing string into said URL; and
18 transmitting said response information, including said URL, to said end user.

1 11- A method as claimed in claim 10 wherein, prior to said determining step, said
2 response information is transmitted from said application through one or more filters.

1 12- A computer program product having computer readable code means of
2 establishing a persistent relationship between an end user device and a server where
3 the server is one of a plurality of servers managed by a dispatcher and the end user
4 device accesses the server using a universal resource locator (URL), the computer
5 program product comprising:
6 computer readable code means of receiving, at the dispatcher, a request for
7 information from the end user device;

8 computer readable code means of determining, by the dispatcher, which of the
9 plurality of servers to select for satisfying the request;

10 compeer readable code means of creating, at the selected server, a token
11 comprising at least an identifier for the selected server, a date/time stamp, and a key,
12 said key for accessing a storage area for information regarding the persistent
13 relationship and the end user device;

14 computer readable code means of inserting the token into the URL; and,

15 computer readable code means of sending, by the selected server to the client
16 device, a response with the token inserted into the URL.

1 13 - A computer program product as claimed in claim 12 wherein said token is
2 encoded using a modified Base64 encoding.

1 14 - A computer program product as claimed in claim 12 wherein said token has a
2 checksum or hash verification field.

1 15 - A computer program product as claimed in claim 14 wherein said hash is a
2 SHA-1 hash computed over said identifier for said selected server, said date/time
3 stamp, and said key.

1 16 - A computer program product as claimed in claim 14 wherein said checksum or
2 hash is encoded using a modified Base64 encoding.

1 17 - A computer program product as claimed in claim 12 wherein said information
2 regarding the persistent relationship is stored as a cookie on said server.

1 18 - A computer program product having computer readable program code for routing
2 a request by an end user device to a particular one of a plurality of redundant servers
3 residing behind a network dispatching mechanism, said program product comprising:

4 computer readable program code for receiving, at the network dispatching
5 mechanism, a request for information indicated by a uniform resource locator (URL);
6 computer readable program code for determining, at the network dispatching
7 mechanism, if said URL contains a valid routing token;
8 if said URL contains a valid routing token, computer readable program code for
9 determining, at the network dispatching mechanism, if a session binding indicated by
10 said routing token is old;
11 if said URL contains a valid routing token and said routing token is not old,
12 computer readable program code for forwarding, by said network dispatching
13 mechanism, the request, including the URL, to the particular server indicated by said
14 valid routing token;
15 computer readable program code for removing, by said particular server, said
16 valid routing information from the URL;
17 computer readable program code for storing, by said particular server, routing
18 information removed from said valid routing token, where said valid routing information
19 can be accessed subsequently by an outbound data stream filter during the processing
20 of an outbound reply related to said request;
21 computer readable program code for accessing, by said particular server, a
22 storage location where information regarding a session between the particular server
23 and the end user device is stored; and
24 computer readable program code for inserting, by said particular server, said
25 session information into said request.

1 19 - The computer program product as claimed in claim 18 wherein additional filtering
2 of the URL is done prior to the forwarding step.

1 20- A computer program product having computer readable program means of
2 sending information to a requesting end user from an application over a session
3 wherein said application resides at one of a plurality of redundant servers residing
4 behind a network dispatcher, said computer program product comprising:

computer readable programming means of receiving response information from
 said application, said response information including a URL (uniform resource locator);
 computer readable programming means of determining if a key cookie has been
 used for storing session information between said end user and said application;
 if a key cookie has been used for storing session information, computer readable
 programming means of retrieving a session key from said key cookie;
 if a key cookie was not used for storing session information, computer readable
 programming means of retrieving said session key from a control block;
 computer readable programming means of removing all cookies from said
 response information;
 computer readable programming means of storing said removed cookies in a
 predetermined storage area;
 computer readable programming means of updating said URL to indicate the
 removal of said cookies;
 computer readable programming means of creating a sticky routing string;
 computer readable programming means of updating a date/time stamp in said
 sticky routing string;
 computer readable programming means of inserting said sticky routing string into
 said URL; and
 computer readable programming means of transmitting said response
 information, including said URL, to said end user.

21- A computer program product as claimed in claim 20 wherein, prior to said
 determining step, said response information is transmitted from said application through
 one or more filters.

~~22~~ - A network dispatcher for establishing a persistent relationship between an end
 user device and a server where the server is one of a plurality of servers managed by
 said network dispatcher, said network dispatcher comprising:

4 means for receiving a request for information from said end user device, said
5 request for information including a uniform resource locator (URL);

6 means for determining which of the plurality of servers to select for satisfying
7 said request for information;

8 means for creating, at said selected server, a token comprising at least an
9 identifier for the selected server, a date/time stamp, and a key, said key for accessing a
10 storage area for information regarding the persistent relationship and the end user
11 device;

12 means for inserting the token into the URL; and

13 means for sending, by the selected server, a response with the token inserted
14 into the URL to the client device.

1 23 - A network dispatcher as claimed in claim 22 wherein said token is encoded using
2 a modified Base64 encoding.

1 24 - A network dispatcher as claimed in claim 22 wherein said token has a checksum
2 or hash verification field.

1 25 - A network dispatcher as claimed in claim 24 wherein said hash is a SHA-1 hash
2 computed over said identifier for said selected server, said date/time stamp, and said
3 key.

1 26 - A network dispatcher as claimed in claim 24 wherein said checksum or hash is
2 encoded using a modified Base64 encoding.

1 27 - A network dispatcher as claimed in claim 22 wherein said information regarding
2 the persistent relationship is stored as a cookie on said server.

Abstract

A method, system and program product are presented for enabling a session, as defined by a series of related transactions to perform a unit of work, to be created between a client and a particular server where the server is managed by a dispatcher. Modifications to the Uniform Resource Locator (URL) are used to create a method of transferring information from the client to the server. The server implements a server-side storage area (cookie jar) to temporarily store information about the client and the session so that the client is routed to the same server for successive messages in the same session and no reliance is made upon an ability by the client to store or return cookies.




Fig. 1

Unmodified:

201 http://www.ibm.com/sales/index

Modified:

205

203 http://www.ibm.com/X@%3.as3.cx.A24/sales/index




Fig. 2

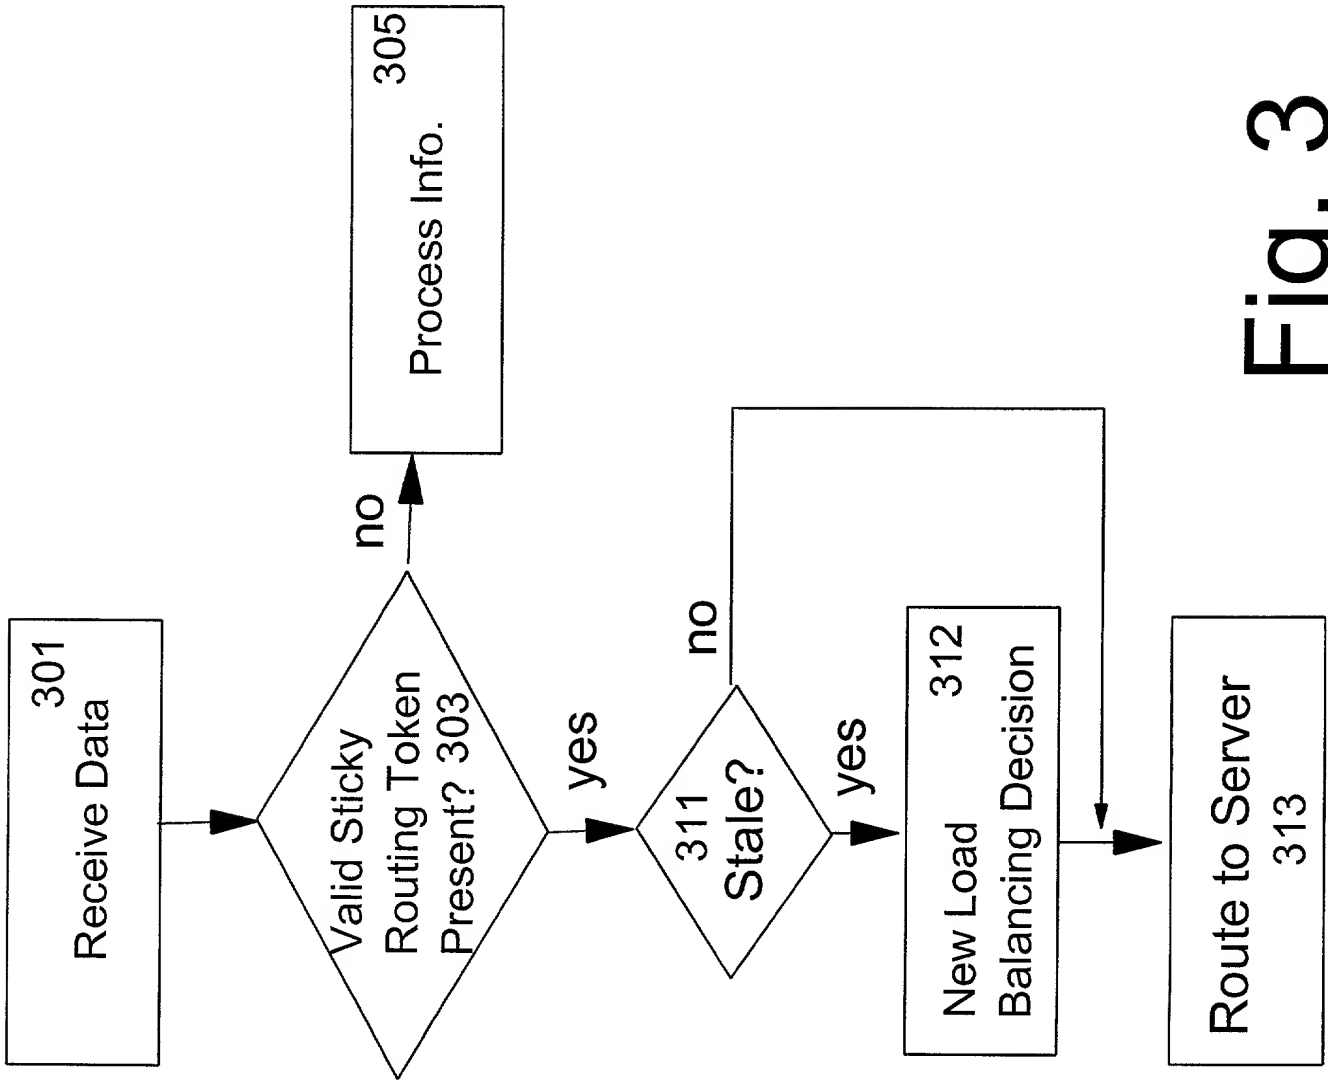


Fig. 3

00000000 8022560

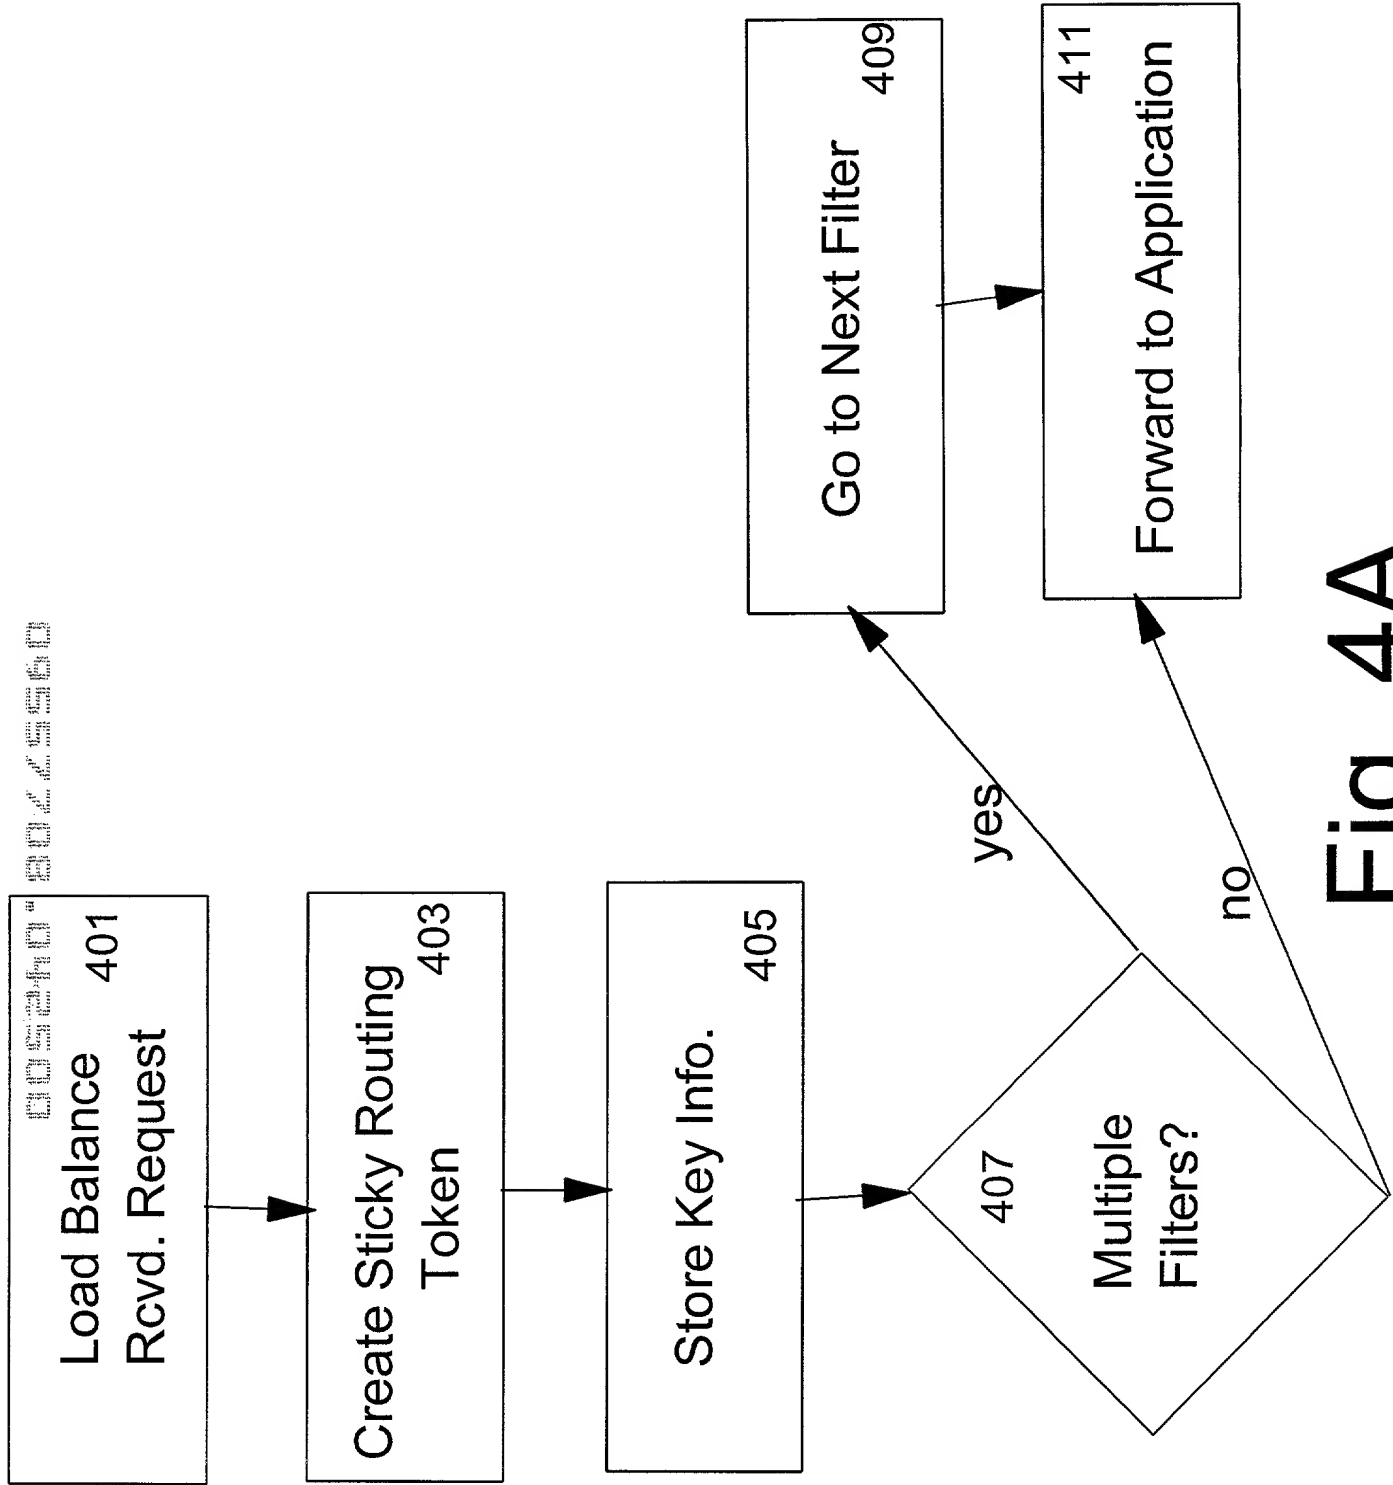


Fig. 4A




Fig. 4B




Fig. 5




Fig. 6

003940" 8022550

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION**

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name; I believe I am an original, first and joint inventor of the subject matter which is claimed and for which a patent is sought on the invention entitled:

URL-Based Sticky Routing Tokens Using a Server-Side Cookie Jar

the specification of which is identified by the attorney (IBM) Docket Number appearing above.

I hereby state that I have reviewed and understand the contents of the above- identified specification, including the claims.

I acknowledge the duty to disclose information which is material to the patentability of this application in accordance with Title 37, Code of Federal Regulations, §1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, §119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

<u>Number</u>	<u>Country</u>	<u>Day/Month/Year</u>	<u>Priority Claimed</u>
---------------	----------------	-----------------------	-------------------------

I hereby claim the benefit (a) under Title 35, United States Code, §119(e) of any U.S. application listed below and identified as a provisional application or (b) under Title 35, United States Code, §120 of any U.S. application listed below and not identified as a provisional application, and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior U.S. application in the manner provided by the first paragraph of Title 35, United States Code, §112, I acknowledge the duty to disclose information material to the patentability of this application as defined in Title 37, Code of Federal Regulations, §1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application

Prior U.S. Applications

<u>Serial No.</u>	<u>Filing Date</u>	<u>Status</u>
-------------------	--------------------	---------------

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

As a named inventor, I hereby appoint the following attorneys and/or agents to prosecute this application and transact all business in the Patent and Trademark Office connected therewith: Bruce A. Clay, Reg. No. 32,121; Gregory M. Doudnikoff, Reg. No. 32,847; Edward H. Duffield, Reg. No. 25,970; Jerry W. Herndon, Reg. No. 27,901; Gerald R. Woods, Reg. No. 24,144; Jeanine S. Ray-Yarletts, Reg. No. 39,808; Joseph C. Redmond, Jr., Reg. No. 18,753; John E. Hoel, Reg. No. 26,279; Christopher A. Hughes, Reg. No. 26,914; and Edward A. Pennington, Reg. No. 32,588;

AND also,

Send all correspondence to: Jeanine S. Ray-Yarletts, IBM Corporation T81/062; PO Box 12195; Research Triangle Park, NC 27709.

First Inventor:

Christopher A. S. Gage

Signature: _____

Date

Residence: 205 Crawley Place, Raleigh, North Carolina 27615

Citizenship: United Kingdom

Post Office Address: same as residence

Second Inventor:

John R. Hind

Signature: _____

Date

Residence: 5005 Barton's Enclave Lane, Raleigh, North Carolina 27613

Citizenship: United States of America

Post Office Address: same as residence

Third Inventor: Marcia L. Peters

Signature: _____ **Date** _____

Residence: 712 Lochgarton Lane, Raleigh, North Carolina 27614

Citizenship: United States of America

Post Office Address: same as residence